



# Cyber Security 2026: What's Really Coming for Your Business

*A predictive look ahead for:*

- *Small Businesses*
- *Enterprises*
- *Government/Edu Institutions*

***iNVI<sub>S</sub>IQ.net***

# *A quick word about the numbers: Before the data police show up.*



If you run a business or lead an organization, you think in **calendar years** and **quarters**: 2024 is done, 2025 is “last year” the minute the fireworks stop, and everyone wants to know what 2026 is going to throw at them.

Cybersecurity doesn’t work like that.

Most of the big, reliable data sources – regulators, law-enforcement reports, large security vendors – only publish **fully checked, final numbers a year or so after the fact**. As of **January 2026**, that means:

- **2024** is the latest year where we have **complete, audited data**
- **2025** numbers are still scattered across partial updates, previews, and “we’ll confirm that later” footnotes

There is no neat “February 1st, all 2025 stats magically appear” moment. The internet is not that organized.

So in this report we’ve done the sensible thing:

- We use **2019–2024** as the **hard data backbone** – the bits that have been counted, argued over, and finalized
- Where there are **early 2025 indicators**, we treat them as **signals, not gospel**
- And we focus on what you actually need for 2026: **direction, scale, and impact on real-world decisions** – not pretending we know every last incident from last year down to the decimal

If you’re looking for a perfect spreadsheet of 2025, you’ll be waiting a while.

If you’re looking for a clear picture of **where the trend is pointing and what that means for SMBs, enterprises, government, and education in 2026**, this is built on the best verified data available today – and we’ll update it as the late-running official numbers catch up.

# Overview

## Total Reported Cybercrime Losses-USA

**\$16.6B**

+33% YEAR OVER YEAR

The FBI's Internet Crime Complaint Center logged **859,532 complaints** and **\$16.6 billion** in reported losses for 2024, a **33% jump** over 2023 – the highest ever recorded

## Published Ransomware Attacks-Global

**5,414**

+11% YEAR OVER YEAR

Global tracking of publicly reported ransomware incidents shows **5,414 organizations hit in 2024**, up **11%** from 2023, with activity spiking in Q4

## Businesses Reporting Attacks

**+40–50%**

HIGHER FOR MEDIUM & LARGE FIRMS

Recent government and industry surveys indicate that **around half of businesses** report at least one cyberattack in the past 12 months, rising to **70%+** for medium and large organizations in some markets.

# Overview (Cont.)

## Average SMB Incident Cost

**\$254,445**

TYPICAL RANGE: \$120K – \$1.24M

Studies focused on small and mid-sized businesses put the **average total cost of a cyber incident** at roughly **\$254,445**, with many cases falling between **\$120,000 and \$1.24 million** – and outliers hitting several million.

## Education & Public-Sector Hit Rate

**60–90%+**

SECTORS UNDER CONSTANT FIRE

A U.S. study found **82% of K-12 schools** had at least one cyber incident between July 2023 and December 2024 – more than **9,300 confirmed incidents across roughly 5,000 schools**.

Put bluntly: if you run a school or a public service and *haven't* seen an incident yet, you're not off the radar – your attacker is just still working through the queue.

## Average Ransomware Downtime

**21–24 DAYS**

WEEKS, NOT HOURS

Ransomware remains the heavyweight champion of disruption: recent reports show **average downtime of roughly 3–4 weeks** per incident, with the cost of unplanned outages for larger organizations measured in **five figures per minute**.

# Narrative

The cyber “market” didn’t so much perform in 2024 as it did **detonate slowly in everyone’s inbox**.

On the numbers side, the FBI’s IC3 report logged **859,000+ complaints** and **\$16.6 billion in losses**, a one-third jump in a single year.

Ransomware crews, apparently unimpressed by global outrage, pushed out **over 5,400 publicly reported attacks**, with the final quarter of 2024 the most active on record.

For businesses, the picture is blunt. Across multiple surveys, roughly **40–50% of firms** now say they’ve taken a hit in the past 12 months, with attack rates climbing higher as you move up the size chart. Medium and large organizations are bigger targets, but that doesn’t mean small ones are safe; it just means they’re quieter when they go down.

Small and mid-sized businesses are absorbing real damage. The average incident cost lands around **\$254,000**, with many cases comfortably into six or seven figures once you add downtime, recovery, legal and lost revenue. [Microsoft Dynamics+2SensCy+2](#)

Ransomware in particular is still the disaster of choice: once it lands, you’re typically counting **weeks of disruption**, not hours, and the meter doesn’t stop when the systems come back online.

Education and the public sector, meanwhile, are learning a harsh lesson in statistics. In the UK alone, **over 90% of universities** and the majority of colleges and secondary schools reported cyber incidents in the last year. [Tom's Hardware](#)

These are not organizations with vast security teams and limitless budgets. They run on trust, email, and aging infrastructure – exactly what modern attackers like best.

As we step into **2026**, this is the backdrop:

- Losses are at record highs.
- Ransomware is still driving long, expensive outages.
- A meaningful slice of SMBs, enterprises, governments, and educational institutions have now experienced at least one serious incident.

The trend is not subtle. The only real question for 2026 is *who* gets hit next, and *how prepared* they are when it happens.



# ***2026 Predictions: What Practically Everyone Will See***

# *Section One - When AI Decides It Wants Your Money*



# ***Prediction #1***

## ***AI-Written Phishing Becomes Just “Normal Email Trouble”***

### **Quick reminder:**

Phishing = messages (usually email) where someone pretends to be someone else to trick you into clicking, paying, or sharing information.

For years, cybercrime was mostly the digital version of junk mail: badly written emails, weird grammar, someone offering you millions if you'd just “kindly assist.”

That era is over.

Attackers have quietly hired **AI** as unpaid staff. Now you get **AI-powered campaigns**, which is a fancy way of saying:

- They **research you first** – scraping your website, LinkedIn, press releases, even your school calendar or council meeting schedule.
- They **write emails and messages that sound like real people** – your CEO, a vendor, a parent, a department head. No broken English, no obvious red flags.
- They **adjust on the fly** – if you ignore the first email, the system tries a slightly different angle, tone, or timing, just like a good salesperson that refuses to go away.
- And when they get access, they can spin out **new malware variants** – new “flavors” of malicious code – fast enough that your security tools struggle to recognize them.

On top of this, we now have **deepfakes** and **synthetic identities**.

- A **deepfake** is audio or video that's been twisted by AI so it looks and sounds like a specific person – your CFO, your superintendent, your mayor – saying things they never actually said.

- A **synthetic identity** is a “person” built mostly out of thin air: some stolen information, some made-up details, AI-generated photos, fake history. They can be used to open accounts, get approved as a vendor, or slip into your systems like a ghost with paperwork.

Put those together and you get this new reality:

You can receive a **perfectly written email**, followed by a **very convincing voice call**, backed by **documents and IDs that pass a quick glance**, and all of it can be fake from top to bottom.

That's what we're walking into in 2026 – not more noise, but more **believable lies**.

Jan 4, 2026 iNVIQ

# ***Impact on SMBs***

For small and mid-sized businesses, this is where it hurts the most.

- Your staff already live in email all day. Now the fake stuff will **look exactly like the real stuff**: same logo, same sign-off, same tone.
- “Can you just pay this invoice today?” will look like the hundreds of other real requests they get, except this one quietly drains your account.
- Because many SMBs don’t have strict approval workflows, a **single well-timed AI email** can jump straight from inbox to bank without anyone blinking.

End result: more “we thought it was legit” stories, fewer funny scam emails to laugh at, and one very annoyed accountant.

# ***Impact on Enterprises***

Big companies already drown in email. AI just makes the water deeper.

- Attackers will **segment** you like a marketing department: finance gets payment lures, HR gets resume lures, IT gets fake alerts, executives get carefully tailored “board-level” nonsense.
- Internal jargon, project names, even nicknames can be scraped from internal or public sources and woven into the email so it feels familiar.
- Security awareness training that relies on “spot the bad grammar” instantly becomes useless. The emails won’t be illiterate; they’ll be **polite and on-brand**.

You don’t get more noise. You get **more believable noise**.

# ***Impact on Government & Education***

Government agencies and schools are prime targets because they’re busy, overloaded, and run largely on email.

- School admins and city staff will see “**official-looking**” **messages** that mimic real forms, portals, and logos almost perfectly.
- Fake requests for records, password resets, vendor payments, or grant adjustments will slip in right next to real ones.

- In education, staff who are used to acting quickly “for the kids” will be pushed to move fast on what looks like a routine, urgent administrative request.

The old advice of “look for spelling mistakes” becomes about as useful as checking the weather to diagnose a broken server.

Jan 4, 2026 iNVI<sup>SIQ</sup>

# ***Prediction #2***

## ***Deepfake Scams Become a Boring Line in Incident Reports***

Today, a deepfake is headline material. By late 2026, it's just a checkbox:

Channel: email / portal / phone / deepfake voice/video

## ***Impact on SMBs***

Small and mid-sized businesses often lean heavily on trust and relationships.

- A “CFO” calling with a cloned voice asking for a **last-minute payment** before the weekend will feel completely normal to a busy bookkeeper.
- Owners who pride themselves on “knowing their people” will discover that **knowing a voice** is now a liability, not a strength.
- Internal fraud and external fraud will start to look very similar – the difference is that in one case, the voice belongs to someone you actually know, and in the other it belongs to a laptop.

SMBs will need to learn a painful lesson: **familiar voice is not proof**.

## ***Impact on Enterprises***

Enterprises have more layers, but they also have more surface area.

- Deepfake video or audio will be used to **pressure internal teams**: IT help desks, finance approvers, legal, even security itself.
- “I’m on the road, just do it and I’ll sign later” will be delivered via a perfect replica of the executive’s voice, with a matching email trail.
- Internal control systems that quietly assumed “spoken confirmation from X is sufficient” will start to generate incidents and audit findings.

Security teams will still hate deepfakes. They'll just hate them as a **routine ticket type**, not as a rare science experiment.

Jan 4, 2026 iNVI<sup>SIQ</sup>

# ***Impact on Government & Education***

Government and education live on **phone calls, video meetings, and in-the-moment decisions.**

- Emergency services, school admins, and agency staff are constantly making judgment calls based on who they think they're talking to.
- If voice and video are officially downgraded as proof, they'll need **clear, written rules** about when a call is enough and when additional checks are mandatory.
- Expect more requirements for **recorded approvals, standardized forms, and multi-step verification** for anything involving money, access, or sensitive data.

It will feel bureaucratic at first. Then the first deepfake-based disaster hits somewhere else, and suddenly the extra steps look like common sense.

# *Fringe 2026 Predictions: If These Hit, the Rules Change*

*These aren't "maybe, sort of, could happen."*

*These are: "if just one of these lands on the evening news, everyone's policies get rewritten by Monday."*

## *Prediction #3*

### *A Major Breach Involves a "Coworker" Who Never Existed*

We already know fake employees are a thing. This isn't theory.

- The U.S. Justice Department has been busy tearing apart schemes where **North Korean IT workers** used stolen or fake identities to get remote jobs at **over 100 U.S. companies**, posing as normal staff and walking off with data and money.
- One U.S. cybersecurity firm, **KnowBe4**, publicly admitted they accidentally hired a **North Korean operative** who passed their checks and joined the team like a normal employee until the truth came out.
- Another case: an Arizona woman ran a "**laptop farm**" that let North Korean workers pretend to be U.S.-based remote employees at **300+ companies**, generating about **\$17 million** before the FBI shut it down.

So "fake employee" is not a prediction. It's already on the scoreboard.

The next step—the fringe piece—is this:

A large, well-known organization discovers that a **key "employee" or vendor contact never existed as a real human being**. Not a stolen identity. A mostly fabricated one, stitched together with AI.

A large, well-known organization discovers that a **key “employee” or vendor contact never existed as a real human being**. Not a stolen identity. A mostly fabricated one, stitched together with AI.

Think of the recipe:

- AI generates a **photo** that looks like a real person.
- AI writes a **resume**, creates **social media history**, and spins up **references**.
- Deepfake tech handles any **video interviews**.
- The identity slides through remote hiring or vendor onboarding, picks up access to systems and money, and quietly does its job... until something breaks.

Jan 4, 2026 iNVIQ

# Impact on SMBs

SMBs will get hit by this in the most painful way possible: through “**helpful experts**” and **contractors**.

- A “remote IT contractor” or “cloud consultant” shows up with clean paperwork and a friendly email style.
- They never come on-site. Everything is email, ticketing systems, and the occasional video call—which AI can happily fake.
- They’re granted admin passwords, VPN access, or keys to your financial systems because “we need this person to get things done.”

When the breach hits, everyone insists they *know* this person. They’ve “worked” with them for months. The moment you realize they do not exist outside of pixels and documents is the moment you rethink every shortcut you ever took with onboarding.

# Impact on Enterprises

Enterprises are built for this kind of mistake.

- HR and vendor teams are drowning in volume. If the resume looks fine and the background check doesn’t scream, the candidate moves forward.
- A synthetic identity can land in a **mid-level role**—developer, analyst, vendor support—somewhere with enough access to cause real damage but not enough visibility to attract early scrutiny.
- When the breach is finally traced back and investigators say, “We can’t find any evidence that this person exists in the real world,” the fallout won’t stop at IT.

You’ll get:

- Audit reports that use phrases like “**systemic failure of identity assurance**.”
- Board-level questions: “How many other fake people did we hire?”
- Regulators asking whether you violated sanctions or other laws by paying a foreign adversary who was pretending to be “Kyle from Denver.”

After one big public case, **identity proofing** (verifying that a person is actually real, not just “has documents”) goes from “nice extra” to “non-negotiable.”

# *Impact on Government & Education*

For government and education, a synthetic “employee” isn’t just embarrassing—it’s existential.

- A fake “consultant,” “researcher,” or “temporary staffer” with fabricated credentials could get direct access to **citizen records, student data, financial systems, or even sensitive infrastructure**.
- Background checks that only look at documents and databases won’t save you from an identity that was built specifically to satisfy those databases.
- When the story hits the press, the headline won’t be “advanced synthetic identity fraud.” It’ll be “Government hired a ghost with access to everything.”

That’s how you end up with:

- Congressional or parliamentary hearings
- Immediate crackdowns on remote hiring and vendor onboarding
- New regulations tying funding or licenses to **strong identity proofing** for anyone with elevated access

One high-profile synthetic employee breach is enough to make every public-sector HR and procurement team suddenly very interested in how they know a person is real.

# ***Prediction #4 (Fringe): Voice and Video Get Demoted From “Proof” to “Maybe”***

For decades, big decisions have been made on the basis of:

- “I recognize his voice.”
- “We saw her on video.”
- “He called me personally and told me to do it.”

Deepfakes are about to make all of that look adorably naive.

We already have:

- Cases where **executive deepfakes in video calls** trick employees into moving large sums of money. [Citi+1](#)
- Governments and regulators warning that deepfake technology is being used to target finance, hiring, and high-trust processes. [Monetary Authority of Singapore+1](#)

The fringe prediction is that by the end of 2026, we see **at least one major legal, regulatory, or insurance decision** that effectively says:

“Voice and video alone are not reliable evidence of who approved what.”

That’s the moment the old “I heard it from the boss” defense goes in the trash.

## ***Impact on SMBs***

SMBs do an enormous amount of business on verbal trust:

- Owners approve large payments on the phone.
- Vendors get “quick green lights” because “we’ve known them for years.”
- Staff are trained informally: if it *sounds* like the owner or the long-time supplier, you move.

If insurers and banks start drawing a hard line—

as in: **“You relied only on a phone call? That’s on you.”**—the entire culture has to shift.

That means:

- “Put it in writing” stops being annoying and starts being **policy**.
- Big transfers need at least one system-based approval, not just “yeah, I said yes on the phone.”
- Emergency changes must go through **some** secondary check, even if it’s just a callback to a known number or a code phrase known only to your team.

You’ll hear a lot of grumbling at first. And then you’ll hear one SMB’s story about losing everything to a cloned voice, and suddenly the grumbling quiets down.

Jan 4, 2026 iNVI<sup>SIQ</sup>

# ***Impact on Enterprises***

Enterprises already have procedures, but people routinely bypass them “just this once” when someone high up is on the line.

- A senior executive calls and says, “I’m in transit, just push it through”—and people do.
- Help desks reset access for “important people” who are “having a nightmare at the airport.”
- Procurement fast-tracks exceptions because “the VP called personally.”

Once there’s a public case where a court or regulator says, “You should have known a voice could be faked,” those shortcuts turn into career risks.

Expect:

- Policy language like: “Verbal requests from any individual, regardless of role, must be confirmed via approved channels for high-risk actions.”
- Internal audits that specifically hunt for **voice-only approvals** and flag them as findings.
- Training that basically says: “If someone is yelling on the phone for you to bypass the system, the answer is no.”

In other words: the system becomes the source of truth, not the voice.

# ***Impact on Government & Education***

Government and education live on phones and video calls. That’s how work gets done.

- Schools constantly handle “urgent” calls about pickups, access, and student records.
- Public-sector workers get calls from people claiming to be officials, partners, vendors, or journalists.
- In a crisis, decisions are often made in minutes, not hours, with a lot of trust placed in who seems to be on the other end.

If voice and video are officially downgraded as proof, they’ll need:

- Clear **written rules** about when a call is enough and when additional verification is mandatory.
- More use of **standard forms, secure portals, and multi-step approvals** for anything involving money, access, or sensitive information.
- Better logging so they can prove later who actually approved what.

It will feel like bureaucracy dialed up to eleven.

But if one deepfake case ends up affecting public safety or student safety, the extra steps will immediately look cheap.

Jan 4, 2026 iNVIQ

These two fringe predictions sit on top of things that are **already happening**:

- Fake workers hired into real companies under stolen or doctored identities.
- Deepfake voices and faces used in interviews, job scams, and executive impersonation. [2](#)

We're not guessing from thin air. We're just following the curve and asking a simple question:

“What happens when criminals stop testing this on the edges and go straight for the biggest, most visible targets—**and it works?**”

# *Section Two – Ransomware Economics: SMB & Education as Prime Meat*





Jan 4, 2026 iNVI<sup>SIQ</sup>

# *Ransomware used to be described as a “once-in-a-career disaster”.*

**Now it’s starting to look more like “line item on the budget you pretend doesn’t exist.”**

Quick translation before we go on:

**Ransomware** = someone breaks into your systems, **locks your files**, and then demands money to unlock them (or not leak them, or not leak them *again*).

## **What the Numbers Actually Say (Without the Marketing Sugar)**

Let’s start with who’s getting chewed on.

- In 2025, analysis shared by **The Hacker News** says about **70.5% of data breaches** hit **small and medium businesses**, not the big-name giants. [LinkedIn](#)  
Attackers basically looked at hardened enterprises, shrugged, and went, “Fine, we’ll just rob the strip mall instead.”
- Other SMB-focused reports show that between **40% and 72% of small businesses** reported some kind of cyber incident in 2024–2025, with ransomware and phishing leading the pack. [DeepStrike+1](#)
- Education is not “a bit targeted.” It’s getting **steamrolled**: one 2025 analysis found educational institutions were getting hit with an average of **4,388 cyberattacks per organization, per week** in Q2 2025 – more than double the global average. [DeepStrike](#)
- Ransomware against schools, colleges, and universities rose **23% year over year** in just the **first half of 2025**, with **average ransom demands around \$556,000** per incident. That’s just the demand, not the total recovery bill. [Cybersecurity Dive+2](#) [Higher Ed Dive+2](#)

So yes, the “prime meat” metaphor is unfortunately accurate:

- **SMBs**: lots of them, often under-protected, usually busy, sometimes underinsured.

- **Education:** tons of data, old systems, no spare budget, and a moral obligation to get back online fast.

From the attacker's point of view, that's not a victim profile. That's a **business opportunity**.

Jan 4, 2026 iNVI<sup>SIQ</sup>

# 2026 Predictions: Things You'll See Whether You Like It or Not

## Prediction #1 – Ransomware Becomes a Recurring Expense, Not a Freak Accident

We're heading toward a world where:

“We were hit by ransomware once”  
becomes  
“We get hit, disrupted, or extorted every few years unless we deliberately change something.”

Why? Because:

- The **tooling is industrial now**. Ransomware is sold as a service (“RaaS”) where one group builds the malware and other groups rent it. It’s basically franchising for criminals.  
[DeepStrike](#)
- Attackers don’t need million-dollar paydays. They’re happy with **many small and medium payouts**, especially from SMBs and schools that just want to get back to work.  
[DeepStrike+1](#)

So instead of thinking “one in a million,” unprepared organizations should be thinking “**cyber hurricane season**”:

- You may not get flattened every year.
- But if you live in the wrong place with the wrong roof, you don’t get to act surprised when it happens again.

## How This Lands by Segment

### SMBs

- Ransomware becomes **part of the risk profile**, like theft or equipment failure.

- You start seeing more businesses that have been hit **twice by different gangs** in a 3–5 year window, because once you pay quickly, your name might as well go on a door marked “fast cash.”
- Insurance renewals get cranky: higher premiums, bigger exclusions, and more questions like “Why is your only backup plan called ‘hope?’”

Jan 4, 2026 iNVIStQ

# Enterprises

- Big companies will still get hit, but they'll increasingly treat it like **containable but expensive maintenance**: downtime, clean-up, external forensics, PR, repeat.
- The real pain moves from “can we recover?” to “how many times are we going to do this before the board loses patience?”
- You'll see more internal conversations that sound like: “We technically survived... but what did this do to our margins, our brand, and our regulator's blood pressure?”

# Government & Education

- For public bodies and schools, ransomware starts to look like a **recurring unfunded mandate**: no one budgets for it, but it keeps showing up.
- Districts and agencies get pushed into **tough choices**: pay ransoms they can't legally admit to, or rebuild systems at huge cost and face public anger.
- Staff and students experience it as “the network is down again” and “we're back to paper forms for a bit,” while behind the scenes the bills pile up.

# *Prediction #2 – Attackers Shift From “Big Game” to “Big Volume”*

We've already seen the trend: **70.5% of breaches hitting SMBs** instead of large enterprises in 2025. [LinkedIn](#)

Translation: the wolves discovered small animals are easier to catch and taste just as good.

In 2026, expect more of this:

- Lower individual ransom demands, **more victims**.
- “We'll encrypt your data” plus “we'll leak your data” plus “we'll harass your customers” – three forms of pressure in one neat, horrible package

# *Segment Impact*

## SMBs

- Attackers don't need \$10 million from you. They're thrilled with **\$50k-\$250k** if they can repeat that hundreds of times a year across different victims.
- Automated scanning tools find exposed services, old VPNs, weak remote desktops, and toss your business onto an attack list without anyone even knowing your company name.
- You start to see more stories that go, "We thought we were too small to be interesting" from companies that clearly were not.

## Enterprises

- Big firms still get targeted, but with **more careful, high-stakes campaigns** while the "bulk ransomware" market focuses on smaller fish.
- Attackers know enterprises are more likely to have backups and negotiators, so they adjust tactics: hit third-party providers, suppliers, and subsidiaries instead.
- The result is **more supply chain ransomware** – your vendors get hit, and you get the downtime.

## Government & Education

- Schools and local government become **volume targets**: many organizations, similar systems, predictable chaos when they're taken offline.
- Because they're public-facing and politically sensitive, attackers know they can generate **faster pressure** on decision-makers to pay, or at least to stay quiet and "make it go away."

# *Fringe 2026 Predictions: If These Happen, the Economics Flip*

Now for the uncomfortable “this will sound insane until it happens once” part.

## **Prediction #3 – Ransomware Gets a “Subscription” Model**

We already have Ransomware-as-a-Service on the criminal side. The next twist is on the victim side:

“Pay us once, and we’ll **put you on a list of companies we promise not to hit again...** and maybe we’ll even tell our friends.”

Is this basically extortion with a loyalty program? Yes.

Does it fit the logic of the criminal market? Also yes.

How it could play out:

- A gang hits you, you pay, you get an **offer**: pay a bit more for “ongoing protection” where they *claim* they won’t target you again.
- They might even sell that “protection” data to other gangs: “Don’t bother with this one, they’re under our ‘umbrella.’”
- Some victims, especially SMBs with no faith in their own security, quietly say yes.

## **Segment Impact**

### **SMBs**

- The worst-case outcome: some SMBs quietly treat this like **buying safety from the mob** because they don’t see a realistic security alternative.
- For them, it becomes a **shadow subscription**: backup service, antivirus, cloud apps... and one “do-not-attack” deal they will never put in writing.

# Enterprises

- Large companies are less likely to cross this line openly—they have auditors and regulators—but you might see **gray-area deals** via third parties, “data deletion” fees, or extended “support” contracts with the gang that hit them.
- If it emerges that a household-name enterprise has effectively paid for a “don’t hit us again” plan, the regulatory fallout will be spectacular.

## Government & Education

- If a public body is ever caught paying for an unofficial “subscription” not to be attacked, the political damage will make the technical damage look small.
- The more likely variant is: attackers **claim** they’ll leave public institutions alone for a price, but still hit them anyway, because there is no court to enforce the contract in criminal-land.

Either way, this shifts ransomware from a **one-off crime** to an ongoing “service relationship,” which is as disturbing as it sounds.

# ***Prediction #4 – Ransomware Costs Start Showing Up in Ratings and Funding***

Right now, most of the pain is internal: downtime, bills, the odd news story, then back to normal.

Fringe prediction: in 2026, **ransomware track records** start showing up where it really hurts—in **credit ratings, insurance conditions, and funding decisions**.

What this could look like:

- A major ratings agency or insurer quietly starts weighting “number of serious ransomware incidents” when evaluating a business or sector.
- Bond markets or lenders adjust terms for organizations with a history of paying ransoms or long outages.
- Education and public bodies see ransomware profile show up in **grant criteria, oversight reviews, or state audits**.

## ***Segment Impact***

### **SMBs**

- A pattern of “we keep getting wiped out for a week” starts to affect how banks and insurers treat you.
- Premiums go up, deductibles go up, and you get more “we’d like to see your controls before we renew this policy.”
- Eventually, “we just wing it on security” starts to sound less like scrappy entrepreneurship and more like **bad risk management** to everyone with money at stake.

### **Enterprises**

- For big firms, ransomware history could become part of how analysts judge **operational resilience**.
- Too many incidents, too much downtime, and suddenly investors and regulators start questioning leadership, not just the firewall.
- You’re no longer just explaining a breach to the press; you’re explaining it to rating agencies who can make your cost of borrowing worse for years.

# Government & Education

- School districts, universities, and agencies may see ransomware response show up in **oversight and funding**:
  - “You’ve had X major incidents, what did you actually fix?”
  - “Why are you still running systems that go down for weeks at a time?”
- A district with repeated, badly handled incidents may find itself under pressure from state or federal bodies—not just to recover, but to **justify why they’re still allowed to run critical systems that keep failing**.

If ransomware economics ever make that jump—from **IT problem** to **credit and funding problem**—the conversation changes overnight. You’re no longer patching servers; you’re defending your ability to get money.

Put bluntly:

- **SMBs and education** aren’t just “in the mix”; they’re the main course.
- The economics of ransomware are shifting toward **volume, recurrence, and leverage**.
- 2026 is the year organizations either start treating that as a structural risk... or keep pretending it’s a freak accident until the next “freak accident” knocks them offline.

# ***Section Three – Cloud & SaaS Concentration Risk: The Internet Monoculture***



# Once upon a time, every company had its own dusty little server room.

Once upon a time, every company had its own dusty little server room.

Terrible, but at least when it broke, **only you** suffered.

Now we've gone all-in on **cloud** and **SaaS**:

- **Cloud** = renting someone else's computers over the internet.
- **SaaS (Software as a Service)** = renting someone else's software over the internet.

- **Cloud** = renting someone else's computers over the internet.
- **SaaS (Software as a Service)** = renting someone else's software over the internet.

It seemed clever: "Let the big boys handle it."

Now we've discovered what happens when **everyone** lets the same three or four big boys handle it.

## What the Numbers Say (And Why It Should Make You Nervous)

- Around **94% of companies worldwide** use some form of cloud in their operations.
- In the core "infrastructure cloud" market (the stuff under everything else), the **top five providers control about 82% of the market**. Amazon alone has roughly **38%**, Microsoft around **24%**, with Google and a couple of others fighting over the rest.

So the internet is now basically:

"A very large pile of businesses sitting on the same small handful of platforms, hoping nothing important breaks on Tuesday."

We've already had a preview of what that looks like when something does go wrong:

- In July 2024, a **faulty CrowdStrike update** pushed to Windows systems triggered what's been called the **largest IT outage in history**: around **8.5 million devices** crashed, disrupting airlines, banks, retailers, hospitals, and governments worldwide, with estimated losses of **at least \$10 billion**.
- In 2024 and 2025, multiple **Microsoft 365** outages knocked out Outlook, Teams, and related services globally, forcing businesses to discover what "no email, no calendar, no docs" really feels like in the middle of the workday.
- Analysis of major cloud outages from 2018–2025 keeps showing the same names: **AWS, Azure/Office 365, Google Cloud, Cloudflare, CrowdStrike** – the core plumbing everyone is leaning on.

Meanwhile, regulators and risk folks are starting to use phrases like “**systemic risk**” and “**concentration risk**” – basically, “we might have built a giant, fragile tower out of the same blocks.”

Jan 4, 2026 iNVI<sup>SIQ</sup>

# Prediction #1 – Big Cloud/SaaS Outages Keep Happening, But With More Collateral Damage

We're not going back to everyone running their own mail server in a broom closet.

Cloud usage is still growing, and the **big providers are getting bigger**, not smaller. So the safe bet is:

**Major multi-hour or multi-day outages at one of the big providers happen again in 2026**, and every time they do, more of the economy stalls at once.

We've seen the template:

- One bad update or network change in a large provider →
- Millions of devices or users impacted →
- Flights grounded, payments delayed, support lines jammed →
- A lot of executives discovering their “resilience plan” was mostly positive thinking.

## Impact on SMBs

For SMBs, “the cloud” often translates to **one SaaS stack for everything**:

- Email, files, CRM, accounting, chat, HR – all in one platform, from one vendor.
- When that platform goes down, your entire business goes into **forced meditation**:
  - No quotes, no invoices, no bookings, no support tickets.

In 2026, it's safe to assume:

- More SMBs will experience **full workdays lost** because “our provider had an issue.”
- A minority will have backup workflows (alternate email, local copies of key documents); most will not.
- Some will discover that **basic things like how to contact customers** are completely locked behind a single vendor's login page.

You wanted to “simplify IT.” You succeeded. You simplified your failure modes too.

## Impact on Enterprises

Enterprises love cloud because it scales, and they hate it because when it breaks, **it scales the outage too**.

- A single outage at a major cloud or SaaS provider now means **thousands of internal systems** can be disrupted at once.
- Internal dependencies (identity, SSO, logging, monitoring) often **all live in the same cloud**. So when that goes down, you lose not just the app, but your ability to see what's

happening.

In 2026:

- More big firms will experience **cross-organizational outages** where multiple regions, departments, and brands go dark at the same time.
- “It’s just email” will be revealed as nonsense, because that same identity platform is tied to **VPNs, dev tools, financial systems, and even physical access in some buildings**.
- Boards and regulators will ask, *“And why exactly is every critical function tied to the same vendor, again?*

Jan 4, 2026 iNVIQ

# Impact on Government & Education

For government agencies and schools, cloud concentration looks like this:

- One identity provider for staff and students.
- One learning platform.
- One productivity suite.
- One finance/HR system.

When that one vendor goes offline:

- Classes stop.
- Forms don't submit.
- Payroll gets "interesting."
- Public communication stalls right when you need it.

In 2026, expect:

- **More news cycles about "a single outage" disrupting multiple public services across a region or country at once.**
- More parents discovering that if the learning platform or email is down, **the school basically can't talk to them.**
- More governments quietly realizing they've put essential public functions onto the same two or three platforms everyone else uses – and that outages don't care who's a "critical service."

# ***Prediction #2 – Single-Vendor Dependency Becomes Embarrassingly Obvious***

“Concentration risk” is a nice formal term. It means:



**You've put way too many eggs in one basket, then balanced that basket on a greased pole in a windstorm.**

By 2026, you'll see a lot more organizations realizing, usually the hard way, that:

- **One identity provider** controls who can log in to anything.
- **One SaaS platform** controls both internal work and communication with customers.
- **One region in one cloud provider** controls an entire business line.

## **Impact on SMBs**

For SMBs, the penny drops like this:

- There's a cloud/SaaS outage.
- Staff stare at each other for a bit.
- Someone says, “Can we at least send a text blast to customers?”
- Everyone realizes all contact info is inside the dead system.

In 2026, more SMBs will:

- Discover they have **no offline copy** of key customer and vendor data.
- Realize that if their main SaaS provider is hacked or locked, **they have no plan B** beyond “wait and hope.”
- Start asking uncomfortable questions like, “If our v

## **Impact on Enterprises**

Enterprises already know about concentration risk; they just don't always act on it.

- Risk teams produce decks about “systemic cloud dependencies.”
- Architects nod.

- Everyone goes back to building more stuff on the same platform because **it's convenient and the contract is already signed.**

Jan 4, 2026 iNVIQ

By 2026, after a few more large incidents:

- More enterprises will be forced into **actual diversification**:
  - backup identity paths,
  - cross-cloud failover for at least some critical workloads,
  - on-prem or alternate providers for absolutely-can't-fail functions. [orx.org+1](#)
- Internal meetings will move from “We should think about multi-cloud one day” to “We cannot explain to regulators why we depend on one vendor for payments, trading, and customer access.”

## Impact on Government & Education

Public bodies and schools are starting to get lectures from above:

- EU's **DORA** and similar rules are already nudging financial institutions to **avoid heavy dependence on any single tech provider** and consider multi-cloud. [AIinvest+1](#)
- Expect the same logic to creep into **public-sector guidance**: “Don't put everything on one SaaS or cloud and then act surprised when it fails.”

In 2026, that probably looks like:

- More **policy language** about resilience and vendor diversification.
- Early pilot projects where a few critical services are deliberately **hosted or backed up somewhere else**, so one outage doesn't take down the entire agency or district.
- Auditors starting to ask: “If Vendor X goes down for 48 hours, what specifically still works?”

# *Fringe 2026 Predictions: If These Happen, the Whole Conversation Changes*

## **Fringe Prediction #3 – A Cloud/SaaS Outage Is Officially Treated as “Systemic Infrastructure Failure”**

We've already had a test run:

- The 2024 CrowdStrike/Microsoft incident showed what happens when a **single faulty update** hits a widespread client base: airlines, finance, healthcare, retail, government—all scrambling at once, with an estimated **\$10B+ impact**.

Fringe call for 2026:

**One major cloud or SaaS outage triggers a formal systemic-risk response – not just “vendor messed up,” but “this is a national or sector-level infrastructure failure.”**

That could mean:

- A financial regulator or central bank publicly labels a specific cloud provider as “**critical infrastructure**” and demands higher resilience, failover, and reporting standards.
- Governments open formal investigations into **how much of their essential services** depend on a handful of tech vendors they don’t actually control.
- New rules requiring large institutions (banks, exchanges, maybe big healthcare systems) to **prove they can survive a multi-day outage of any single cloud or SaaS provider**.

## **Impact on SMBs**

SMBs won't get these memos directly—but they'll feel the side effects:

- Vendors will start selling “**regulator-approved resilience**” as a feature, with higher costs passed downstream.
- Some SMB-oriented platforms may quietly move to **multi-region or multi-cloud** setups, and then use that as a marketing point: “We stay up when others don’t.”

- A few unlucky SMBs will feature in case studies: “Here’s what happened to our customers when Vendor X went dark.”

Jan 4, 2026 iNVI<sup>SIQ</sup>

# Impact on Enterprises

For enterprises, this prediction means:

- Cloud/SaaS concentration risk moves from being a **slide in a risk presentation** to a **condition for operating in regulated sectors**.
- Boards start asking: “Can we survive 72 hours without Provider A?” and if the answer is “absolutely not,” things get awkward.
- Some firms will be forced to invest in truly independent backup paths:
  - separate identity providers,
  - offline-capable critical systems,
  - contractual guarantees about data portability and emergency failover.

You know it's serious when resilience projects stop being optional “nice-to-have” budget lines and start being **license-to-operate requirements**.

# Impact on Government & Education

Government and education may be dragged along by example:

- If banks and major financial institutions are told “don’t rely on just one provider,” it’s only a matter of time before **hospital systems, power grids, transportation, and education** get the same advice—if not formal mandates.
- Public bodies may be told to build **“service continuity plans”** that assume a major cloud provider is down for multiple days, not just an afternoon.

That’s when “we moved everything to the cloud to save IT costs” turns into “we now need a cloud continuity budget as well.”

# *Fringe Prediction #4 – “We Can Live Without Any One Vendor” Becomes a Selling Point*

Right now, tech marketing loves phrases like:

- “All-in on [Vendor Name]!”
- “Unified on a single platform!”

It sounds tidy. It also sounds like **please punch me here**.

Fringe prediction:

**By the end of 2026, some organizations will brag, publicly, that they can operate for days without any single cloud or SaaS provider, and customers, regulators, and investors will treat that as a serious advantage.**

## **Impact on SMBs**

For most SMBs, this will start small:

- A few will keep **critical contact data, key documents, and processes** accessible outside their primary SaaS provider.
- Some will choose SaaS platforms that make it easy to **export and locally mirror** vital data or run in a degraded mode when the cloud is down.
- Over time, “we can still take orders and talk to customers when X is offline” becomes a quiet but real differentiator.

The rest will keep rolling the dice, and occasionally discover what “totally dead in the water” feels like on a random Thursday.

## **Impact on Enterprises**

Enterprises will turn this into a metric:

- “We can continue core operations for **72 hours** if [Provider] is down.”
- “Customer-facing services are multi-homed across **two clouds and on-prem failover**.”
- “Identity and access can fall back to **an independent provider** if our primary platform fails.”

This becomes part of:

- Investor story (“we’re operationally resilient”),
- Regulator conversations (“we don’t create systemic risk”),
- Customer trust (“we stay up when competitors go dark”).

It’s the resilience version of having a strong balance sheet.

Jan 4, 2026 iNVI<sup>SIQ</sup>

# Impact on Government & Education

For public-sector and education:

- Some early adopters—larger universities, national agencies—will start to **publicly document** their ability to keep essential services running through major vendor outages.
- This might include:
  - alternate communication channels for parents and citizens,
  - offline or locally-hosted fallbacks for critical records,
  - non-cloud-dependent contingency processes for things like payroll and emergency alerts.

You'll know this prediction landed when:

- A big outage hits a major cloud provider,
- Several agencies and school systems go down,
- And one or two step up and say, “We stayed online; here’s how,” and suddenly **everyone** wants a copy of their playbook.

---

Big picture for 2026:

- We’ve built an **internet monoculture** on top of a small number of mega-vendors.
- Outages that used to be annoying are now **economy-sized events**.
- SMBs, enterprises, governments, and schools are all discovering that “move everything to one cloud” solved yesterday’s problems and quietly built tomorrow’s.

The question isn’t whether the monoculture exists—it does.

The question is who, in 2026, decides to stop pretending it’s invincible.

# ***Section Four – Identity-Centric Attacks: When Your Login Is the Attack***



# *Because why break the door when you can just borrow the keys?*

We spent years building higher walls: firewalls, VPNs, fancy boxes with blinking lights. Attackers looked at all that, shrugged, and said:



**“Or we could just log in like everyone else.”**

That's what **identity-centric attacks** are: instead of smashing the system, they **steal or abuse accounts** and stroll through the front door.

Quick translation of the buzzwords:

- **MFA fatigue / push bombing** – spam you with login approval prompts until you finally hit “Approve” just to make it stop. [Hoxhunt+1](#)
- **Token theft / session hijacking** – steal the “ticket” that proves you’re already logged in, so attackers skip the password and MFA entirely. [Push Security+1](#)
- **OAuth abuse** – trick you into clicking “Allow” on one of those “This app would like to access your account” pop-ups, and boom: long-term access without another password prompt.

## **What the Numbers Say (Spoiler: They’re All About Identity)**

This isn’t a niche side quest; it *is* the main story now.

- One 2024–2025 global survey found **93% of organizations** suffered **two or more identity-related breaches** in the past year. Identity isn’t just “a factor”; it’s the battlefield.
- Proofpoint reports that in 2024, **99%** of monitored organizations were targeted for **account takeover**, and **62%** actually had at least one account taken over (average: **12 takeovers per org**).
- Verizon/IBM-backed stats show **79% of web app compromises** involved **breached credentials**, not fancy exploits.
- Infostealer malware stole **around 1.8 billion credentials in 2025**, with stolen passwords and session cookies appearing in **86% of the breaches** studied. Credentials are now industrial-scale raw material.

- Microsoft and others have tracked **hundreds of thousands of MFA fatigue attacks** in a year. About **1% of users** will blindly accept the first push request, and incident responders say **roughly a quarter of recent attacks involve fraudulent MFA prompts**.
- In the U.S., the FBI says **\$262M+** has already been stolen in **account takeover scams in 2025 alone**, across **5,100+ complaints** – and that's just what got reported.

Meanwhile, every second blog from serious shops now says the same thing in slightly different fonts:

**“Identity is the new perimeter.”**

Translation: your usernames, passwords, tokens, and “Sign in with X” buttons are now more important than whatever metal box you’ve put in front of them.

Jan 4, 2026 iNVIQ

# 2026 Predictions: What You'll See Whether You Like It or Not

## Prediction #1 – Most Incidents Start With “Someone Logged In Who Shouldn’t Have”

We used to say “they exploited a vulnerability.”

Increasingly, the story is:

- They **phished someone** or used AI to fake a login page.
- Or they **bought stolen credentials** from an infostealer log.
- Or they **tricked someone into approving an MFA prompt** at 11:30 p.m.
- Or they **stole an active session** and skipped MFA altogether.

By 2026, the safe bet is:

For SMBs, enterprises, government, and schools, **the majority of serious incidents start with a misused identity**, not a zero-day exploit.

## Impact on SMBs

For small and mid-sized businesses, this lands like a brick:

- Your “security stack” might be nothing more than a cloud login and SMS codes.
- If an attacker gets a password plus **one good MFA approval** or **one stolen session cookie**, they’re in. And once they’re in, it all looks legitimate:
  - Email from the real account
  - File access from the usual IP range
  - Activity during normal working hours

The incident report won’t say “they bypassed our perimeter.”

It’ll say “**they logged in as Lisa from accounting and just... stayed.**”

## Impact on Enterprises

Enterprises are already seeing this:

- Identity providers and SSO (single sign-on) systems have become **crown jewels**: compromise one account there and you can pivot into multiple apps.
- Logs show “successful login” from a known device, so legacy alerting doesn’t even flinch.

- Internal tools, cloud consoles, data lakes – all fair game once the identity is trusted.

By 2026, expect:

- More big-name breaches where the root cause is “**valid credentials used in clever ways**,” not a missing patch.
- More board conversations that sound like: “We spent how much on perimeter tools... and they still just signed in?”

Jan 4, 2026 iNVI<sup>SIQ</sup>

# Impact on Government & Education

Public bodies and schools have painfully simple realities:

- Shared devices, shared accounts, and a lot of “**just log in as admin at the front desk.**”
- Staff logging into critical systems from home, on personal laptops, with browsers that haven’t seen an update since the Obama administration.

In 2026, more incidents will look like:

- A student, parent, or staff account compromised and then reused to move sideways into more sensitive systems.
- Tokens and sessions stolen from personal devices that also have games, random browser extensions, and who-knows-what running.

The root cause won’t be “the firewall failed.” It’ll be “we treated identities like spare keys in a bowl.”

## ***Prediction #2 – “We Have MFA” Stops Sounding Impressive***

**MFA = Multi-Factor Authentication** – using something more than just a password (like a push notification, code, or app prompt).

A few years ago, saying “we have MFA” made you sound responsible.

Now we have:

- Attack data showing **79% of BEC (business email compromise) victims in 2024–2025 actually had MFA turned on.**
- At least **25% of recent attacks** involving **MFA push abuse** – attackers spamming login approvals until someone gives in.

So the safe prediction:

**By 2026, nobody serious will treat “we turned on MFA” as proof of anything.**

**The question becomes: what kind of MFA, how is it configured, and how do you detect when it’s being abused?**

# Impact on SMBs

For SMBs:

- You turned on SMS codes or push notifications and thought you were done.
- Then someone gets **MFA-bombed** during a busy day: 10 prompts, 20 prompts, 30 prompts... they finally hit “Approve” just to make the noise stop.
- Or worse, their phone number gets SIM-swapped and all your precious codes go straight to the attacker. SIM-swap fraud jumped **over 1,000% in 2024** in some markets. [Specops Software](#)

By 2026, “We have MFA” becomes:

- “We have **phishing-resistant** MFA for the important stuff,”
- “We limit push prompts,”
- “We actually look at weird login behavior.”

Everyone else is just doing “**MFA theater**.”

# Impact on Enterprises

Enterprises will get grilled on:

- Which users have **strong methods** (hardware keys, FIDO2, passkeys) vs **weak methods** (SMS, endless allow/deny prompts).
- How often they **track MFA abuse** – not just failed logins, but repeated prompts and approvals at odd hours.
- Whether they have **identity threat detection** in place (ITDR) to spot “valid login, very invalid behavior.” [Petronella Technology Group, Inc.+1](#)

The story shifts from “we turned it on” to “**we operate it like it matters**.”

# Impact on Government & Education

Government and education have unique pain:

- You’ll still have older staff, contractors, and community partners stuck on SMS-based MFA because that’s all they can handle.
- You’ll have students and staff bringing **their own devices**, with all the risk that implies.

In 2026, the safe outcome is:

- MFA becomes **tiered**: stronger methods for high-risk roles (finance, admin, IT, records), simpler ones for low-risk access.
- Auditors start asking **how many high-privilege accounts still rely on the weakest possible MFA**.

If the answer is “most of them,” expect some interesting recommendations.

Jan 4, 2026 iNVIStQ

# Fringe 2026 Predictions: If These Land, Identity Rules Get Rewritten

## Fringe Prediction #3 – BYOD and Shared Devices Get Smacked With Real Limits

We all loved **BYOD** – *Bring Your Own Device* – because it was cheaper and staff liked using their own phones and laptops.

Identity attacks are quietly ruining that party.

- Infostealers on personal devices happily slurp up **work passwords, cookies, and tokens** in the background. [DeepStrike+1](#)
- Shared family computers with “just one more Chrome extension” become a buffet for session hijacking.

Fringe call:



**By the end of 2026, at least in higher-risk sectors, BYOD and shared device policies get serious teeth – not just “guidelines,” but hard technical and contractual limits.**

## Impact on SMBs

For SMBs, this will feel like a culture shock:

- Today: “Use whatever laptop and phone you have, just get the job done.”
- Tomorrow in certain roles:
  - “If you handle finance, HR, admin, or IT, you **don’t log in from your kid’s Fortnite box.**”
  - “Here’s a managed device, or here’s the containerized work profile; use that or don’t handle sensitive stuff.”

You’ll still have BYOD, but not for “**keys to the kingdom**” type access.

# Impact on Enterprises

Enterprises are already inching here:

- More **device posture checks** before granting access: Is it updated? Is it managed? Is there an EDR agent?
- More **split policies**: normal staff can use BYOD for low-risk apps; high-privilege roles must use managed endpoints.

In 2026, fringe becomes reality when:

- A major incident is traced directly to a **personal device compromise** (infostealer, token theft),
- Regulators or insurers say, “You let admin-level access from that thing?”
- BYOD gets quietly gutted for admin, developer, and finance roles.

# Impact on Government & Education

Government and schools often run on shared machines:

- Front-desk PCs, lab computers, classroom stations – dozens of people logging into the same boxes all day.
- Students and staff mixing personal browsing, school work, and sometimes admin functions on the same hardware.

If one ugly incident exposes token theft from a widely shared device, expect:

- **Stricter separation**: student devices vs staff devices vs admin devices.
- “No, you cannot approve sensitive changes from the receptionist’s kiosk.”
- Possibly new guidance saying certain roles **must** use managed, dedicated devices—no exceptions.

When identity is the attack, the device suddenly matters again.

# *Fringe Prediction #4 – Identity Telemetry Becomes a Regulated Data Set*

## **When the “boring logs” suddenly matter more than your shiny dashboard**

Right now, most organizations treat **identity telemetry** – logs of who logged in from where, with what device, doing what – as something only security nerds care about.

Here's the fringe prediction:



**By late 2026, at least one major regulator or industry body starts treating identity telemetry as mandatory evidence – something you must keep, protect, and produce to prove you had control over your environment.**

Think:

- Financial regulators asking for **identity event history** after outages or fraud:
  - “Show us when these accounts were accessed, from which locations, and what they did.”
- Education/government bodies asking to see **access patterns** around sensitive record changes.
- Insurers demanding identity telemetry as part of breach investigations: “If you can’t show us who did what, when, why should we write the check?”

## **Impact on SMBs**

For SMBs, this is going to be... educational:

- A lot of small shops have **weak or no logging** beyond what their cloud or SaaS vendor keeps by default.
- If the expectations move to “you need at least basic identity logs and some way to review them,” SMBs will either:
  - Lean heavily on managed providers, or
  - Start paying for logging/monitoring they never budgeted for.

On the plus side, **good identity logging** is the difference between “we have no idea what happened” and “here’s exactly where it went wrong.”

Jan 4, 2026 iNVIStQ

# Impact on Enterprises

Enterprises already have logs coming out of their ears, but identity data is often scattered:

- One set in the IdP, another in the VPN, another in each cloud, another in legacy systems.
- Investigations turn into scavenger hunts.

If regulators and insurers start expecting **coherent identity trails**, not just piles of logs, enterprises will:

- Invest more in **identity threat detection and response (ITDR)**, not just SIEM.
- Treat identity data like **financial records**: something you keep, normalize, and can actually read under pressure.

# Impact on Government & Education

For government and education:

- Identity telemetry becoming important means **audit trails on who accessed what student/citizen record and when** become non-negotiable.
- When there's a controversy ("who changed this record?" "who accessed these files?"), showing clean identity logs becomes the difference between:
  - "We can prove it was misused," and
  - "We honestly have no idea."

That's the moment identity moves from "IT plumbing detail" to "**evidence we must have if something goes wrong**".

---

Big picture for 2026:

- Attackers aren't just hacking systems; **they're hacking logins and habits**.
- MFA is necessary but nowhere near sufficient.
- BYOD and shared devices are quietly becoming the weakest links in the identity chain.
- And the organizations that treat identity telemetry and access patterns as **core infrastructure**, not a side effect, will be the ones that can actually explain – and limit – the damage when someone inevitably "logs in who shouldn't have."

# About This Report

This report was not written by a committee of beige consultants in matching vests.

It's a **2026 cybersecurity outlook** for SMBs, enterprises, government, and education — written in plain English, with a healthy disrespect for buzzwords and marketing fluff.

The goal is simple:

- Show you **where the risk is actually moving**,
- Translate that into real-world impact for the four sectors we care about,
- And do it without putting you to sleep or trying to sell you a magic box that “solves cyber.”

---

## How to Read the Data (Before You Start Yelling at the Numbers)

Cyber data doesn't run on business time.

You close your books on December 31 and move on.

Regulators and big vendors take another 6–18 months to publish their final, cleaned-up numbers.

As of **early 2026**:

- The latest **fully audited, stable data** from most law-enforcement bodies and major vendors is complete **through 2024**.
- **2025** figures are a mix of partial releases, previews, early studies, and “we'll revise this later” charts. Some are solid. Some will age badly.

So this report does the sane thing:

- Uses **2019–2024** as the **hard data backbone** (the years that have actually been counted and argued over).
- Adds **early 2025 indicators** where they're credible and directionally useful.
- Builds a **2026 outlook** on top of that — focusing on **trends and direction**, not pretending we have final 2025 spreadsheets from the future.

If you're looking for exact, final “2025 global everything” statistics in January 2026, you're going to be disappointed in more places than this report.

# AI Disclosure: Who Did What

Short version:

**This report is AI-assisted, but human-led.**

What that actually means:

- **Humans chose the topics.**

Real people decided that AI-driven attacks, ransomware economics, cloud/SaaS monoculture, and identity abuse are the four pillars that matter for 2026. Not a model hallucinating themes at random.

- **Humans did the heavy lifting.**

Structure, judgment, editing, deciding what stays, what gets cut, and what's too ridiculous to print — that's all human work.

- **AI was used as a power tool, not a ghostwriter.**

It helped:

- Pull and cross-check public stats faster than any intern with a browser.
- Draft first-pass language that was then rewritten, sharpened, or deleted.
- Compare multiple sources on the same topic so we weren't relying on one lonely chart.

- **Humans set the tone and the line.**

The sarcasm, the sector focus (SMB / enterprise / gov / edu), and the difference between “safe” vs “fringe but possible” predictions — all human calls.

If you're wondering whether this is “AI generated Slop”: **NO!**

AI **assisted**; humans **directed and approved**. If you hate something in here, blame the people, not the silicon.

# *Sources & Methods (Without the 40-Page Appendix)*

s isn't based on "a guy on social media said so."

The analysis leans on **public, reputable sources**, including:

- **Law-enforcement and government reports**
  - National cybercrime and fraud reports
  - Ransomware and business-email-compromise (BEC) loss figures
  - Sector-specific stats where available (e.g., education, healthcare, public sector)
- **Major vendor and industry research**
  - Annual and quarterly threat reports from large security vendors
  - Cloud and SaaS outage analyses
  - Identity and access-related breach studies (account takeover, MFA abuse, session hijacking)
- **Sector and SMB-focused surveys**
  - Studies on incident frequency, average breach cost, downtime, and ransomware impact for small and mid-sized organizations
  - Education and government exposure to ransomware and service outages
- **News and economic impact coverage**
  - Reporting and post-mortems on large, well-publicized incidents (global outages, major ransomware events, deepfake scams, synthetic identity cases)
  - Independent estimates of business disruption and recovery costs

Where sources disagreed, they were treated like witnesses, not gospel:

- We checked **definitions** (what they mean by "incident," "breach," "attack").
- We looked for **convergence across multiple reports**, not one dramatic outlier.
- We favored **transparent methodology** over "trust us, we're a platform."

This is a **synthesis**, not a single-vendor brochure.

# **What This Is – and What It Isn’t**

This report **is**:

- A **2026 cybersecurity outlook** written for **SMBs, enterprises, government, and education**, not just Fortune 50 security teams.
- A translation of complex threat trends into **plain language and business impact** (What does this mean for how we operate? For budgets? For risk?).
- Opinionated on where things are heading – and honest about what we know vs what we’re inferring.

This report is **not**:

- A “buy this tool and you’re safe” sales deck.
- A deep technical manual for specialists who already read 200-page NIST documents for fun.
- A fantasy document pretending we have perfect, final, globally agreed 2025 data in early 2026.

If all you take away is this:

Attackers already have a plan for the next 12–18 months.

This is our attempt to make sure you’re not the only one at the table who doesn’t.

...then the report has done its job.

# About iNVISIQ

iNVISIQ is not “yet another security vendor with a dashboard fetish.”

iNVISIQ is a **behavior-based cybersecurity company** built around one simple idea:

Most breaches don’t start with elite zero-days.

They start with **predictable human behavior** that attackers know how to exploit.

We focus on four worlds that keep getting punched in the face:

- Small and midsize businesses
- Larger enterprises that still have SMB habits in key areas
- Government and public-sector organizations
- Education: K-12, colleges, universities

Where most players obsess over **tools**, iNVISIQ obsesses over **behavior plus exposure**:

- How your people actually work when they’re tired, overloaded, or under pressure
- How your **vendors, cloud tools, and SaaS platforms** quietly become your biggest attack surface
- How attackers chain those two realities together to get exactly what they want

The result is an ecosystem of tools and services designed to answer one question, over and over:

“Given how your people and your stack really behave,  
what’s most likely to hurt you next – and what do you do about it *before* it does?”

# ***The iNVI~~S~~IQ Toolkit (Where SMOKE Fits In)***

iNVI~~S~~IQ isn't a single product; it's a growing set of tools that all share the same behavioral spine.

**SMOKE** is one of those tools – the flagship, not the whole fleet.

- **SMOKE** is the **external early-warning engine**.

It watches the **outside world** for signals that your vendors, SaaS tools, and other “touch points” are on fire:  
breaches, leaks, critical vulnerabilities, ransomware hits, ugly incident disclosures, and all the fun stuff attackers read before breakfast.

SMOKE's job is to:

- Map the services and providers your business actually depends on
- Track **public risk signals** around those touch points
- Turn that into **plain-English alerts** your IT people and decision-makers can act on, quickly

Around SMOKE, iNVI~~S~~IQ is building and evolving **other modules** to cover:

- Internal behavior and exposure
- Response coaching and decision support
- Ongoing risk storytelling your board can actually understand

Different tools, same philosophy:

- **Behavior first**
- **External + internal reality, not theory**
- **Actionable, not theatrical**

SMOKE is the part that watches the horizon.

The rest of iNVI~~S~~IQ is about making sure you do something intelligent with what it sees.

# ***About the Author – Bradford Allen***



This report is led by **Bradford Allen**, founder of iNVISIQ and the one responsible for the predictions you're about to hold against us later.

The headline credential – not a side note:

**Bradford holds a degree in Applied Behavioral Science (ABS) from National-Louis University.**

That matters for one very specific reason:

Cybersecurity isn't just about networks and boxes.

It's about **people**:

- How they make decisions
- How they cut corners
- How they react when something looks urgent, confusing, or just slightly off

An ABS degree means Bradford's foundation is **rigorous study of human behavior and systems**, not just "I like gadgets."

National-Louis University looked at his work on behavior and said, essentially, "*Yes, this*

Everything in this report flows from that lens:

- Threats are evaluated not just on what's technically possible, but on **how people are likely to respond**.
- Predictions are built around **behavior patterns attackers already exploit**, not marketing buzz.
- SMBs, enterprises, government, and education are treated as **different behavioral environments**, not one generic "user."

Beyond the degree:

- Bradford has **decades of real-world experience** in teaching, coaching, small business, and high-stakes sales — all domains where misunderstanding human behavior costs you real money.
- He's worked in **OSINT and cyber-adjacent roles**, tracking scams, fraud, and breaches from the outside, where the data is messy and late, not neatly packaged.
- He built iNVISIQ specifically to bridge the gap between **what attackers are actually doing** and **how real organizations make decisions under pressure**.

So when this report says, "Here's what 2026 is likely to throw at SMBs, enterprises, government, and education," that isn't coming from a purely technical echo chamber.

It's coming from an **ABS-trained view of behavior** combined with a platform that spends its time staring at how attackers really move.



**NATIONAL  
LOUIS  
UNIVERSITY**

*person is qualified to analyze and predict  
how humans actually act.”*

Jan 4, 2026 iNVIStQ

# Why iNVI~~S~~IQ Is Qualified to Make These Predictions

You don't need another "year ahead" PDF that just rephrases vendor marketing slides.

iNVI~~S~~IQ's predictions come from three overlapping vantage points:

## 1. Behavioral Science (ABS)

- Understanding how humans reliably behave under stress, ambiguity, and overload
- Mapping attacks to **behavioral patterns**, not just CVE numbers

## 2. Real External Signals (via tools like SMOKE)

- Watching which vendors, SaaS platforms, and sectors are actually getting hit
- Tracking how fast – or slow – the news of those hits reaches SMBs, schools, agencies, and enterprises
- Seeing where the same types of failures keep repeating

## 3. Operational Reality in SMB, Enterprise, Gov, and Edu

- Knowing that a 15-person shop, a 3,000-person enterprise, and a school district do **not** think, buy, or respond the same way
- Building predictions that reflect those differences instead of pretending "one size fits all"

Put all that together and you get:

- Forecasts that start with "**Given how people actually behave, what will attackers try next?**"
- Not just "What's the shiny new exploit?" but "**Where will this cause real damage first?**"
- And a view of 2026 that's grounded in **behavior, exposure, and real-world constraints**, not fantasy budgets and perfect compliance.

If you want predictions from someone whose only credential is a pile of certs and a love of blinking lights, this isn't that.

If you want predictions from someone with an **Applied Behavioral Science degree**, running a **behavior-based cybersecurity company** whose tools (including SMOKE) watch what attackers and vendors are doing to **businesses, government, and education right now** – that's exactly what you're reading.

Thank you for taking your time to read this manual.

